Shared and Private VAEs with Generative Replay for Continual Learning

Subhankar Ghosh
Indian Institute of Science
Bengaluru, India
subhankarg@alum.iisc.ac.in

Abstract

Continual learning tries to learn new tasks without forgetting previously learned ones. In reality, most of the existing artificial neural network(ANN) models fail, while humans do the same by remembering previous works throughout their life. Although simply storing all past data can alleviate the problem, it needs large memory and often infeasible in real-world applications where last data access is limited. We hypothesize that the model that learns to solve each task continually has some task-specific properties and some task-invariant characteristics. We propose a hybrid continual learning model that is more suitable in real case scenarios to address the issues that has a task-invariant shared variational autoencoder and T task-specific variational autoencoders. Our model combines generative replay and architectural growth to prevent catastrophic forgetting. We show our hybrid model effectively avoids forgetting and achieves state-of-the-art results on visual continual learning benchmarks such as MNIST, Permuted MNIST(QMNIST), CIFAR100, and miniImageNet datasets. We discuss results on a few more datasets, such as SVHN, Fashion-MNIST, EMNIST, and CIFAR10. Our code is available at https://github.com/DVAEsCL/DVAEsCL.

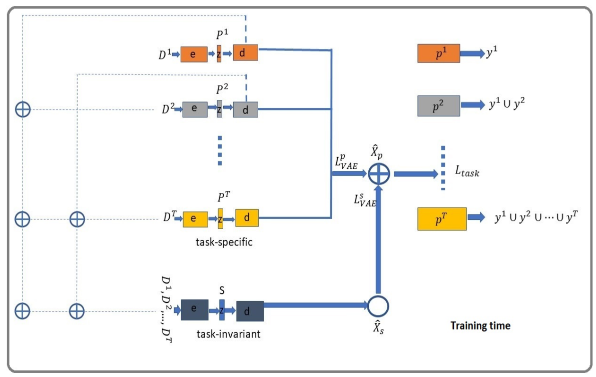


Figure 1. is our model at training time. Architecture growth occurs at the arrival of tth task by adding a task-specific VAE denoted as P^t, and a task-specific ANN denoted as p^t. To prevent forgetting, private VAEs are stored for each task. The first private VAE gets trained using real data from the first task. The second private VAE is trained using real data from the second task and synthesized data corresponding to the first task's classes generated from the first private decoder. Similarly, the tth private VAE sees real data from the tth task, synthesized data from the previous private decoders corresponding to their tasks' classes during training. A shared VAE that is less prone to forgetting yet is also retrained with a small number of generative replay. The plus(+) sign indicates concatenation.

1. Introduction

Humans and animals are capable of continually learning and updating knowledge throughout their lifetime. The ability to accommodate new experiences while retaining previously known knowledge helps to build reusable artificial intelligent systems. Current ANNs achieve impressive performance on many machine learning problems like image classification, object detection, and natural language processing but fail to remember previous knowledge due to a phenomenon called catastrophic forgetting[28] when trained for new tasks. We want our artificial learning agents to have the ability to solve many tasks sequentially under different conditions by developing task-invariant and task-specific skills that enable them to adapt while avoiding forgetting using generative replay quickly. Several approaches have been proposed over the years to alleviate the catastrophic forgetting. The first approach involves dynamically increasing the network's capacity to learn new tasks[44, 34]. The second approach uses regularizers that make the number of parameters remain constant during sequential learning[1, 17, 22, 23, 46]. However, these approaches are not feasible if each task needs a large mem-

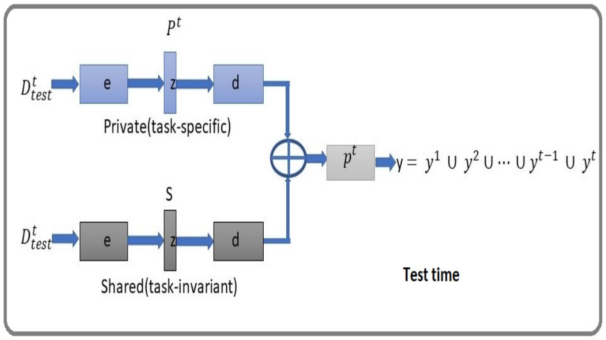


Figure 2. is our model during test time at t^{th} task. The model gets data corresponding to all classes till the tth task and predicts those classes

ory. Another line of approach is to rely either on experience replay[4, 24, 31] or generative replay[33, 37, 40] by storing real data from previous tasks or train generative models, respectively.

We propose a novel continual learning method in which the network is dynamic, uses generative replay, and grows in size with each task. A disjoint space representation composes task-invariant or shared space(fixed size) trained for all tasks, and task-specific or private space grows with each task. Our approach is motivated by the fact that our human brain structure is complex and contains billions of neurons[12]. We may need to move towards a more complex neural network structure in the coming future to make artificial systems solve the works humans can do nowadays. Continual learning can be applied to various practical situations involving privacy issues. The main contributions of these work are summarized as follows:

- We develop a structure-based and generative replay-based model using (T + 1) numbers of conditional variational autoencoders. T is the number of tasks the model has to solve.

- We show the private and shared modules' training is different from[11] though that paper inspires our network's architecture.

- We present results for some datasets and show our model achieves state-of-the-art performance on all the datasets.

2. Related Work

2.1. Continual Learning

The existing continual learning approaches can be divided into architecture-based, regularization-based, and rehearsal-based strategies.

Architecture-based Methods

The first approach to prevent catastrophic forgetting does modification in the network's structure by growing a module for each task either physically or logically[44, 23]. These methods attempt to localize inference to a subset of the network such as columns[26], neorons[10, 45], a mask over parameters[27, 36]. The performance of learned tasks is preserved by storing the learned modules while accommodating new tasks by augmenting the network with new modules. PNNs[35] statically grow the architecture, are immune to forgetting, and utilize prior knowledge via lateral connections to previously learned features. Reference[45] proposed a dynamically expandable network(DEN) that can dynamically decide its network's capacity as it gets trained on a sequence of tasks. These methods impose a computational cost in continual learning where many tasks need to be learned, and fixed capacity memory can not be considered.

Regularization-based Methods

The second family of this field is based on regularization. They estimate the importance of a network's parameters and penalize those weights while switching from one task to another task. There are many existing approaches for penalizing the weights. One of the methods is the elastic weight consolidation[EWC][18], where important parameters have the highest in terms of the Fisher information matrix. In reference[46], the weights are computed online and kept track of how much the loss changes due to change in specific weights and accommodate this information during training. Reference[1] focuses on the change on the activation instead of considering the loss's change. This way, parameter importance is learned in an unsupervised manner. Despite the success gained by these methods, they are often limited by the number of tasks.

Rehearsal-based Methods

The final family of methods of this domain to mitigate forgetting is rehearsal-based. Existing approaches use two strategies: either store few samples per class from previous tasks or train a generative model like GAN[38] or VAE[14] or both to sample synthetic data from previously learned distributions. The iCaRL[31] stores a subset of real data(exemplars). For a constant memory budget, the number of data stored per learned class decreases as the number of tasks increases, so the models' performance decay. Reference[13] proposed two losses called: the less-forget constraint and inter-class separation to prevent forgetting. The less-forget loss minimizes the cosine distance between the features extracted from the old and new models. References[41, 3] introduce a bias-correction layer to correct the original fully-connected layer's output to address the data imbalance between the old and new cate-

gories. A recent study on tiny episodic memories in continual learning are GEM[25], A-GEM[21], MER[32], and ER-RES[5].

The second strategy in this family does not store any data but generates synthetic data using generative models. Reference[37] used generative replay with an unconditional GAN, where an auxiliary classifier needs to determine which classes the generated samples belong. Reference[40] is an improved version of[37], where they used class-conditional GAN to synthesize data. Reference[15] used a generative autoencoder for replay. Synthetic data for previous tasks are generated based on the mean and covariance matrix using the encoder's class statistics. The major limitation of these approaches is the assumption of a Gaussian distribution of the data.

2.2. Space Factorization

In machine learning, multi-view learning is more effective, more promising, and has a better generalization ability than single-view learning[43]. The approaches to tackle multi-view learning aim at either maximizing the mutual agreement on different views of the data or focus on gaining sub-space shared by many views by assuming that the input views are synthesized from that subspace using clustering[6], Gaussian processes[39], etc. So, the concept of factorizing the space into shared and private sub-spaces has been explored[9]. In this paper, we factorize the data space into two parts: shared and private.

3. Shared and Private VAEs with Generative Replay for Continual Learning

We study the problem of learning a sequence of T data distributions denoted as $D = \{D^1, D^2, ..., D^T\}$, where $D^t = \{(X_i^t, Y_i^t, T_i^t)_{i=1}^{n_t}\}$ is the data distribution for the task t with n_t sample tuples of input($X^t \in \mathcal{X}$), target label ($Y^t \in \mathcal{Y}$), and task label($T^t \in \mathcal{T}$). The goal is to learn a sequential function, $f_\theta : D^t \rightarrow \hat{\mathcal{Y}}^t$, for each task, where $\hat{\mathcal{Y}}^t$ are the predicted labels corresponding to t^{th} task. $f_\theta \in (f_S \cup f_P \cup f_p)$, where $f_S : D^t \rightarrow \hat{\mathcal{X}}_S$, $f_P : D^t \rightarrow \hat{\mathcal{X}}_P$, and $f_p : \hat{\mathcal{X}}_S \cup \hat{\mathcal{X}}_P \rightarrow \hat{\mathcal{Y}}^t$. We try to achieve our goal by training two separate modules: private and shared, to mitigate forgetting of prior knowledge. The model prevents catastrophic forgetting in shared and private spaces separately and begins learning f_θ^t where $\theta \in (\theta_S, \theta_P, \theta_p)$ as mapping function from $\mathcal{D}^t$ to $\mathcal{Y}^t$. We use some n samples per class to be synthesized prior to t^{th} task and accumulate the generated data to the current task(t^{th}) to train the model. During training the model with t^{th} task:

$$X^t \leftarrow X^t \cup \hat{\mathcal{X}}_p^{1:(t-1)}$$

The cross-entropy loss function for the f_θ^t mapping corresponds to:

$$L_{task=t}(f_\theta^t, D^t) = - \mathop{\mathbb{E}}_{(\mathcal{X}^t, \mathcal{Y}^t) \sim D^t} \left[\sum_{c=1}^{C} \mathbb{1}_{(c=\mathcal{Y}^t)} log(\sigma(f_\theta^t(\mathcal{X}^t, \mathcal{Y}^t))) \right] \quad (1)$$

Where σ is the softmax function, in learning a sequence of tasks, an ideal f_θ^t maps the input images X^t to their predicted labels $\hat{\mathcal{Y}}^t$.

Variational autoencoders(VAEs)

Autoencoders can effectively learn input space and representation[47, 2]. A VAE is a generative model that follows an encoder-latent vector-decode architecture of classical autoencoder, which places a prior distribution on the input space and uses an expected lower bound to optimize the learned posterior. Conditional VAE is an improved version of the VAE, where data are fed to network with class properties such as labels. The variational autoencoderis a fundamental building block of our approach. Variational distribution tries to find a true conditional probability distribution over the latent space z through minimizing their distance using a variational lower bound limit. The loss function for a VAE is:

$$L_{VAE} = \mathop{\mathbb{E}}_{q_\phi(z|x)} [log(p_\theta(x|z))] - D_{KL}(q_\phi(z|x) \parallel p_{\theta(z)}) \quad (2)$$

Where the first one is the reconstruction loss, and the second term is the KL divergence between $q(z|x)$ and p(z). The encoder predicts μ and $\sum$ such that $q_\phi(z|x) = \mathcal{N}(\mu, \sum)$, from which a latent vector is synthesized via reparametrization process.

The final objective function of our approach for the tth task is:

$$L^{(t)} = \lambda_1 L_{task} + \lambda_2 L_{VAE}^S + \lambda_3 L_{VAE}^P \quad (3)$$

Where, λ_1, λ_2, and λ_3 are regularizer constants to control the effect of each loss component. The working algorithm of this model is presented in Algorithm 1.

3.1. Avoid forgetting

Catastrophic forgetting occurs because of the imbalance in the data between previous and new classes that creates a bias in the network towards the current ones during training, and models almost forget previous knowledge. One of our approach's insights is to decouple the single space learned for all tasks continually into two parts: shared and private sub-spaces. Another approach is the generative replay from previous private modules that concatenated into the current

task's data during training the model with the current task to avoid forgetting. The first private module sees real data of the first task; the second private VAE sees real data of the second task and synthetic data of the first task. Similarly, the third private module gets real data of the third task and synthetic data of the first and second tasks generated from the first and second private decoders, respectively, during training the model with the third task. Whatever data individual private module sees during training, the only shared module gets trained by them. It goes like this till the Tth task.

3.2. Evaluation Matrices

We estimate the resulting model on all previous tasks similar to[25, 8] after training for each new task. We use ACC as the average test classification accuracy across all classes for continual learning to measure our model's performance. To measure forgetting, we calculate backward transfer, BWT that calculates how much learning new tasks have influenced previous tasks' performance. While $BWT < 0$ indicates catastrophic forgetting and $BWT > 0$, learning new tasks has helped improve performance on previous tasks.

$$BWT = \frac{1}{T-1}\sum_{t=1}^{T-1}[R_{T,t} - R_{t,t}] \tag{4}$$

$$ACC = \frac{1}{T}\sum_{t=1}^{T} R_{t,t} \tag{5}$$

ACC is the mean classification accuracy across all tasks. Where $R_{j,i}$ is the test classification accuracy on task i after sequentially finishing learning the jth task.

4. Experiments

This section consists of the datasets and baselines we used in our experiments and the implementation details.

Datasets

We perform our approach on the commonly used benchmark datasets for T-Split continual learning, where the entire dataset is divided into T tasks. We use **5-Split MNIST** and **Permuted MNIST(QMNIST)**[20] previously used in[9, 8, 30, 48], **20-Split miniImageNet**[45] used in[5, 9, 49], **20-Split CIFAR100**[19] used in[25, 21, 9, 48]. We also perform our experiments on **5-Split SVHN**[29], **5-Split CIFAR10**[19], **5-Split Fashion-MNSIT**[42] and **13-Split not-MNIST(EMNIST)**[7]. The datasets' statistics are given in Table 1. We have not used any data augmentation techniques.

Algorithm 1 Continual Learning

Input: $(\mathcal{X}, \mathcal{Y}) \sim D^{all}$
Parameters: $\theta \in (\theta_S \cup \theta_P \cup \theta_p)$
Output: $\mathcal{Y}$

```
D_gen ← {}
for t ← 1 to T do
    for e ← 1 to epochs do
        Compute L_task using (X^t, Y^t) ∈ D^t
        Compute L^S_VAE for the shared module using (X^t, Y^t) ∈ D^t
        Compute L^P_VAE for the t^th private module using (X^t, Y^t) ∈ D^t
        L^(t) = λ1 L_task + λ2 L^S_VAE + λ3 L^P_VAE
        θ' ← θ − α∇L^(t)
    end for
    accuracy ← function TEST(X^t_test, Y^t_test, t)
    for c ← 1 to C do
        C is the replay classes.
        for i ← 1 to n do
            n is the number of samples to be generated per class for the experience replay.
            X̂^i_p ∼ D^{1:t}_gen
        end for
    end for
    X^{t+1} ← X^{t+1} ∪ X̂_p
end for
```

Dataset	#Classes	#Tasks	Input Size	#Train Data(k)	#Test Data(k)
MNIST	10	5	$1 \times 28 \times 28$	50	10
QMNIST	10	5	$1 \times 28 \times 28$	60	50
EMNIST	26	13	$1 \times 28 \times 28$	100	12
F-MNIST	10	5	$1 \times 28 \times 28$	60	10
CIFAR10	10	5	$3 \times 32 \times 32$	50	10
CIFAR100	100	20	$3 \times 32 \times 32$	50	10
SVHN	10	5	$3 \times 32 \times 32$	61	12
mImageNet	100	20	$3 \times 84 \times 84$	50	10

Table 1. Statistics of the datasets. Where, EMNIST = not-MNIST, QMNIST = Permuted MNIST, F-MNIST = Fashio-MNIST, and mImageNet = miniImageNet.

Baselines

We compare with state-of-the-art approaches, including elastic weight consolidation(EWC)[18], Progressive neural networks(PNNs)[35], Hard Attention Mask(HAT)[36], and ACL[9] using implementations given by[9] unless otherwise stated. We compare a few memory-based methods A-GEM[21], GEM[25], ER-RES[5] for MNIST, permuted MNIST(QMNIST), 20-Split CIFAR100, and 20-Split mini-ImageNet. We depend on the implementation provided by[9]. On Permuted MNIST results for SI[48] are taken from[36], for VCL[30], those are taken from[9], and for

Dataset	Encoder	Decoder	Preceptron
MNIST	4 CL	4 DCL	4 CL and 1 FC
QMNIST	4 CL	4 DCL	4 CL and 1 FC
EMNIST	4 CL	4 DCL	4 CL and 1 FC
Fashion-MNIST	4 CL	4 DCL	4 CL and 1 FC
CIFAR10	4 CL	4 DCL	2 CL and 2 FC
CIFAR100	4 CL	4 DCL	2 CL and 1 FC
SVHN	4 CL	4 DCL	2 CL and 2 FC
miniImageNet	5 CL	5 DCL	4 CL and 1 FC

Table 2. The information of the networks used in experiments(CL = Convolution layers, DCL = Deconvolution layers, FC = Fully connected layers).

uncertainty-based CL in Bayesian framework(UCB)[8] are taken from the actual paper.

Dataset	z dimension	#parameters
MNIST	108	199019
QMNIST	108	199019
EMNIST	108	562939
Fashion-MNIST	108	199019
CIFAR10	192	3388428
CIFAR100	192	13790903
SVHN	192	3388428
miniImageNet	96	4107152

Table 3. The second column gives the dimension of latent space for shared and private VAEs for all datasets. The third column shows the number of parameters required to train each dataset.

Implementation Details

The information of the networks used in our model is given in Table 2 for each dataset. For a dataset, the architectures of shared and private modules are the same. We take PyTorch as our working framework. We train each of all datasets for 50 epochs and evaluate our model's performance at the 25th and 50th epochs. We perform experiments using 4, 20, and 100 synthetic samples per class during continual training. The Adam optimizer[16] has been used for all experiments, and the learning rate for the model is 0.0001. The dimension of the latent variables z and the number of parameters used for each dataset are given in Table 3. We take $\lambda_1 = \lambda_2 = \lambda_3 = 1$ at the loss function. We use a Tesla V100 gpu in all our experiments.

5. Results and Discussion

In the first set of experiments, we measure ACC, BWT, and the memory used by our method and compare it against state-of-the-art methods on 20-Split miniImageNet, 20-Split CIFAR100, 5-Split Permuted MNIST, and 5-Split MNIST. Next, we demonstrate the experiments on sequentially learning single datasets such as SVHN, Fashion-MNIST, EMNIST, and CIFAR10.

5.1. Performance on 20-Split miniImageNet Dataset

We divided the miniImageNet dataset into twenty tasks with five classes for each task. We compare our results with several baselines in Table 4. HAT[36] as a regularization-based method with no replay data achieves ACC = 59.45. A-GEM[21] and ER-RES[5] use architecture with 25.6M parameters along with memory replay. They store thirteen images of size $84 \times 84 \times 3$ per class during continual training. Reference[9] used both architecture-based and memory-based approaches together and outperformed other algorithms, and it achieves ACC = 62.07. Our model beats all existing models at a large margin and earns ACC = 100 when we take 100 synthetic samples per class and train it for 50 epochs. The model gives ACC = 77.4 when we take only four generated data per class and train it for 25 epochs. Our model takes 2900 seconds to learn the 20-Split miniImageNet data when trained for 50 epochs and used four samples generated per class. The ACCs for each task are given in Figure 3 using a different combination between the number of epochs the model gets trained and the number of samples used as a generative replay.

Method	ACC(%)	BWT(%)	Arch(MB)	M Replay	G Replay
HAT*[36]	59.45	-0.04	123.6	-	-
PNN**[35]	58.96	0.00	588	-	-
ER-RES*[5]	57.32	-11.34	102.6	✓	-
A-GEM*[21]	52.43	-15.23	102.6	✓	-
ORD-FT[9]	28.76	-64.23	37.6	-	-
ACL[9]	62.07	0	113.1	✓	-
ours	100(0.00)	6.6(0.2)	16	-	✓

Table 4. Results on 20-Split miniImageNet data measuring ACC(%), BWT(%), and memory size(MB) for architecture. M Replay: memory replay or actual data stored and G replay: Generated synthetic data. BWT positive means good. (*) denotes result is reproduced by[9].(**) denotes result is obtained using the re-implementation setup by[36]. All results are averaged over 3 runs and standard deviation is given in parentheses.

5.2. Performance on 20-Split CIFAR100 Dataset

We split the whole dataset into 20 tasks, where each contains five classes. We compare our results with other methods in Table 5. ACL[9] is the most competitive baseline, takes 24.5 MB to save its architecture, takes 13 images per class(1300 images of size $(32 \times 32 \times 3)$ in total) that require 16 MB of memory and achieves ACC = 78.08. HAT[36]

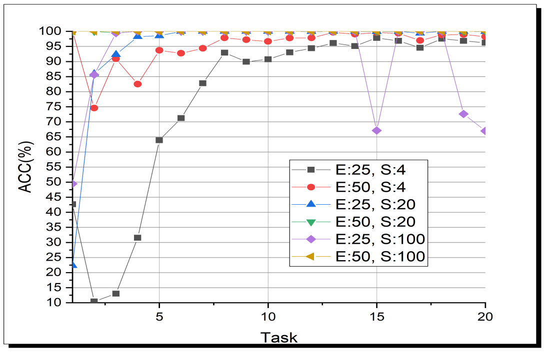


Figure 3. Results for the 20-Split miniImageNet data. E = Number of epochs the model gets trained, S = Number of synthetic data used per class as a generative replay.

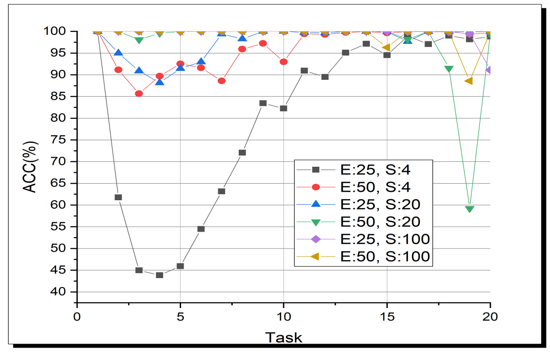


Figure 4. Results for the 20-Split CIFAR100 data. E = Number of epochs the model gets trained, S = Number of synthetic data used per class as a generative replay.

Method	ACC(%)	BWT(%)	Arch(MB)	M Replay	G Replay
HAT*[36]	76.96	0.01	27.2	-	-
PNN**[35]	75.25	0.00	93.51	-	-
ER-RES^o[5]	54.38	-21.99	25.4	✓	-
A-GEM^o[21]	66.78	-15.09	25.4	✓	-
ORD-FT[9]	34.71	-48.56	27.2	-	-
ACL[9]	78.08	0	25.1	✓	-
ours	99.52(0.59)	8.5(0.5)	53	-	✓

Table 5. Results on 20-Split CIFAR100 data measuring ACC(%), BWT(%), and memory size(MB) for architecture. M Replay: memory replay or actual data stored and G replay: Generated synthetic data. BWT positive means good. (*) denotes result is obtained by[9] using original provided code. (**) denotes result is obtained using the re-implementation setup by[36]. (^{o}) denotes result is reported by[5]. All results are averaged over 3 runs and standard deviation is given in parentheses.

does not depend on replay example but needs 27.2 MB for its architecture in which learns task-based attention mask reaching ACC = 76.96. PNN[35] guarantees zero forgetting. Our model outperforms all the existing models in a large margin and achieves ACC = 99.52 when we train the model for 50 epochs and use 100 synthetic data per class. The model gives ACC = 80.6 if we train it for 25 epochs with only four synthetic samples per class. Our model learns 20-Split CIFAR100 data in 1540 seconds if trained for 50 epochs and used only four synthetic samples per class as a generative replay. The ACCs for each task are given in Figure 4 using a different combination between the number of epochs the model gets trained and the number of samples used as a generative replay.

5.3. Performance on Permuted MNIST Dataset

Method	ACC(%)	BWT(%)	Arch(MB)
EWC^{oo}[18]	88.2	-	1.1
HAT^o[36]	97.4	-	2.8
UCB^o[8]	91.44(0.04)	-0.38(0.02)	2.2
VCL*[30]	88.80(0.23)	-7.90(0.23)	1.1
VCL-C*[30]	95.79(0.10)	-1.38(0.12)	1.1
PNN**[35]	93.5(0.07)	zero	N/A
ORD-FT[9]	44.91(6.61)	-53.69(1.91)	1.1
ACL[9]	98.03(0.01)	-0.01(0.01)	2.4
ours	100(0.00)	0.00(0.00)	0.8

Table 6. Results on Permuted MNIST data measuring ACC(%), BWT(%), and memory size(MB) for architecture. BWT positive means good. (^{oo}) denotes result is reported by[36]. (^{o}) denotes result is reported by original work. (*) denotes result is obtained by[9] using original code. (**) denotes result is reported by[5]. All results are averaged over 3 runs and standard deviation is given in parentheses.

One of the popular variants of the MNIST dataset in continual learning is Permuted MNIST. We divide the dataset into five tasks where each contains two classes. We compare our approach with other methods in Table 6. HAT[36] achieves ACC = 91.6 using an architecture of size 1.1 MB. Vanilla VCL[30] improves ACC and BWT by 7 and 6.5, respectively, using k-means core-set-memory size of 200 samples per task(6.3 MB) and architecture size of 1.1 MB. PNN[35] achieves ACC = 93.5 with zero forgetting. ACL[9] achieves ACC = 98.03 takes 0.2 MB for adding 55k parameters for each task, and occupies a total of 2.5 MB memory. Our model outperforms all methods and achieves ACC = 100 when we train it for 50 epochs with four synthetic samples per class as a generative replay and occupies 0.8 MB for its architecture. If we train the model for 25 epochs with only four synthetic samples per class, it gives

ACC = 98.12. The model learns Permuted MNIST dataset in 1300 seconds when trained for 50 epochs with same number of synthetic data. The ACCs for each task are given in Figure 5 using a different combination between the number of epochs the model gets trained and the number of samples used as a generative replay.

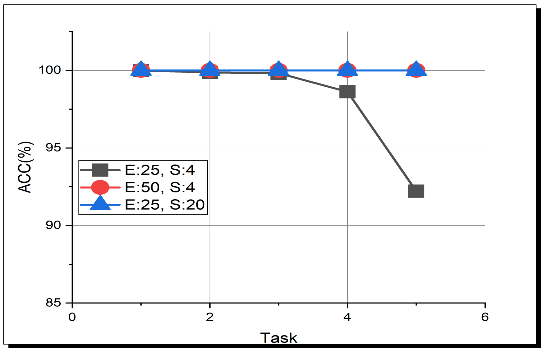


Figure 5. Results for the Permuted MNIST data. E = Number of epochs the model gets trained, S = Number of synthetic data used per class as a generative replay. For other((E:25, S:100), (E:50, S:20), and (E:50, S:100)) combinations ACC = 100% for all tasks, so we did not plot.

5.4. Performance on 5-Split MNIST Dataset

Method	ACC(%)	BWT(%)	Arch(MB)
EWC^{oo}[18]	95.78(0.35)	-4.2(0.21)	1.1
HAT^{oo}[36]	99.59(0.01)	0.00(0.04)	1.1
UCB^{o}[8]	99.63(0.02)	0.00(0.00)	2.2
VCL*[30]	95.97(1.03)	-4.62(1.28)	1.1
VCL-C*[30]	93.6(0.20)	-3.10(0.20)	1.7
GEM*[25]	94.34(0.82)	-2.01(0.05)	6.5
ORD-FT[9]	65.96(3.53)	-40.15(4.27)	1.1
ACL[9]	99.76(0.03)	0.01(0.01)	1.6
ours	100(0.00)	0.00(0.00)	0.8

Table 7. Results on 5-Split MNIST data measuring ACC(%), BWT(%), and memory size(MB) for architecture. BWT positive means good. (^{oo}) denotes result is reported by [8]. (^{o}) denotes result is taken from original work. (*) denotes result is obtained by [9] using original provided code. All results are averaged over 3 runs and standard deviation is given in parentheses.

We divide the MNIST dataset into five tasks, where each consists of two classes. We compare our results with other existing models in Table 7. EWC, HAT, UCB, and Vanilla VCL are regularization-based methods with no memory replay are provided in that Table. Methods relying on memory only(GEM) and VCL with k-means core-set(VCL-C) where 40 samples are stored per task. ACL gives ACC = 99.76 with zero-forgetting outperforming UCB with ACC = 99.63 which uses 40% more memory. ACL uses only architecture growth(no experience replay), where 54.3k private parameters are added for each task resulting in a memory requirement of 1.6 MB for all private modules. ACL's architecture has a total of 420.1k parameters. Our method outperforms all existing models, achieves ACC = 100 with BWT = 0 when we train it for 50 epochs with only four synthetic data per class as a generative replay, and learns the MNIST dataset in 660 seconds. The ACCs for each task are given in Figure 6 using a different combination between the number of epochs the model gets trained and the number of samples used as a generative replay.

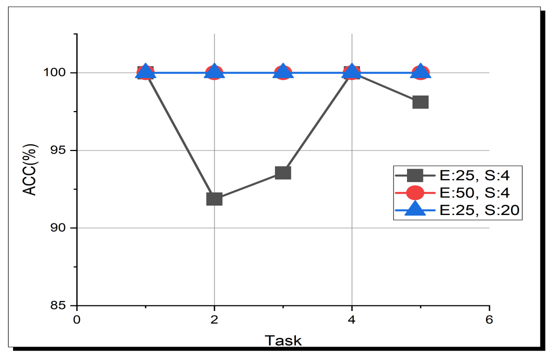


Figure 6. Results for the MNIST data. E = Number of epochs the model gets trained, S = Number of synthetic data used per class as a generative replay. For other((E:25, S:100), (E:50, S:20), and (E:50, S:100)) combinations ACC = 100% for all tasks, so we did not plot.

5.5. Performance on SVHN Dataset

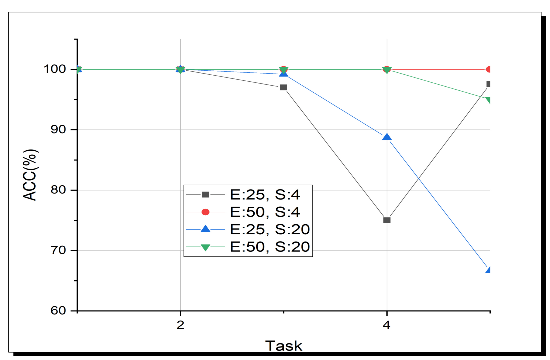


Figure 7. Results for the SVHN data. E = Number of epochs the model gets trained, S = Number of synthetic data used per class as a generative replay.

We divide the SVHN dataset into five tasks, where each consists of two classes. Our model learns the whole dataset continually around 800 seconds when trained for 50 epochs and used only four synthetic samples per class as a genera-

tive replay. The network has a total of 3388428 parameters that occupies 13 MB of memory. The model gives ACC = 100. How the model performs in each task is given in Figure 7 with few combinations between the number of epochs the model gets trained and the number of samples used as a generative replay.

5.6. Performance on Fashion-MNIST dataset

We divide the Fashion-MNIST dataset into five tasks, where each consists of two classes. Our model gives ACC = 100, learns the whole dataset continually around 600 seconds when trained for 50 epochs, and used only four synthetic samples per class as a generative replay. The network has a total of 199019 parameters that occupies 0.8 MB of memory. How the model performs in each task is given in Figure 8 with few combinations between the number of epochs the model gets trained and the number of samples used as a generative replay.

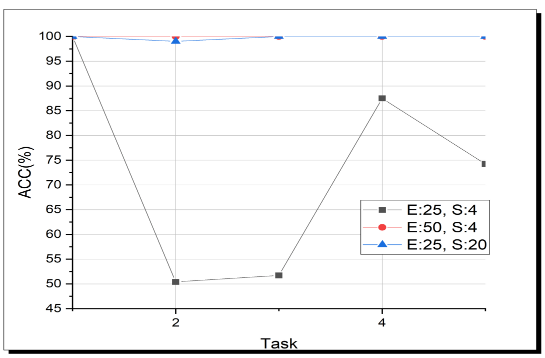


Figure 8. Results for the Fashion-MNIST data. E = Number of epochs the model gets trained, S = Number of synthetic data used per class as a generative replay.

5.7. Performance on EMNIST Dataset

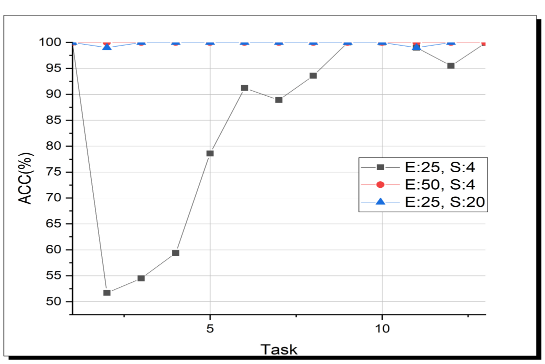


Figure 9. Results for the EMNIST data. E = Number of epochs the model gets trained, S = Number of synthetic data used per class as a generative replay.

We divide the EMNIST dataset into thirteen tasks, where each consists of two classes. Our model gives ACC = 100, learns the whole dataset continually around 1800 seconds when trained for 50 epochs, and used only four synthetic samples per class as a generative replay. The network has a total of 562939 parameters that occupies 2 MB of memory. The model's performance in each task is given in Figure 9 with few combinations between the number of epochs the model gets trained and the number of samples used as a generative replay.

5.8. Performance on CIFAR10 Dataset

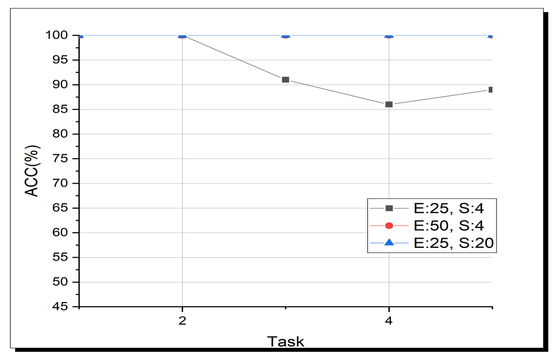


Figure 10. Results for the CIFAR10 data. E = Number of epochs the model gets trained, S = Number of synthetic data used per class as a generative replay.

We divide the CIFAR10 dataset into five tasks, where each consists of two classes. Our model gives ACC = 100, learns the whole dataset continually around 800 seconds when trained for 50 epochs, and used only four synthetic samples per class as a generative replay. The network has a total of 3388428 parameters that occupies 13 MB of memory. The model's performance in each task is given in Figure 10 with few combinations between the number of epochs the model gets trained and the number of samples used as a generative replay.

6. Conclusion

In this paper, we propose a novel hybrid continual learning algorithm that grows in size with each task and factorizes the representation learned for a sequence of tasks into task-invariant and task-specific sub-spaces. The learning method combines generative replay and architecture-based approaches together. We show that it gives excellent performance on datasets having more classes or more tasks like miniImagenet and CIFAR100. We establish a new state-of-the-art on continual learning benchmark datasets. For future work, we are interested in extending this work in class incremental learning.

References

[1] Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 139–154, 2018. 1, 2

[2] Samaneh Azadi, Deepak Pathak, Sayna Ebrahimi, and Trevor Darrell. Compositional gan: Learning image-conditional binary composition. *International Journal of Computer Vision*, 128(10):2570–2585, 2020. 3

[3] Eden Belouadah and Adrian Popescu. Il2m: Class incremental learning with dual memory. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 583–592, 2019. 2

[4] Arslan Chaudhry, Puneet K Dokania, Thalaiyasingam Ajanthan, and Philip HS Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 532–547, 2018. 2

[5] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc'Aurelio Ranzato. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486*, 2019. 3, 4, 5, 6

[6] Kamalika Chaudhuri, Sham M Kakade, Karen Livescu, and Karthik Sridharan. Multi-view clustering via canonical correlation analysis. In *Proceedings of the 26th annual international conference on machine learning*, pages 129–136, 2009. 3

[7] Gregory Cohen, Saeed Afshar, Jonathan Tapson, and Andre Van Schaik. Emnist: Extending mnist to handwritten letters. In *2017 International Joint Conference on Neural Networks (IJCNN)*, pages 2921–2926. IEEE, 2017. 4

[8] Sayna Ebrahimi. Continual learning with neural networks. 2020. 4, 5, 6, 7

[9] Sayna Ebrahimi, Franziska Meier, Roberto Calandra, Trevor Darrell, and Marcus Rohrbach. Adversarial continual learning. *arXiv preprint arXiv:2003.09553*, 1, 2020. 3, 4, 5, 6, 7

[10] Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A Rusu, Alexander Pritzel, and Daan Wierstra. Pathnet: Evolution channels gradient descent in super neural networks. *arXiv preprint arXiv:1701.08734*, 2017. 2

[11] Subhankar Ghosh. Adversarial training of variational auto-encoders for continual zero-shot learning. *arXiv preprint arXiv:2102.03778*, 2021. 2

[12] S Herculano-Houzel. The human brain in numbers: a linearly scaled-up primate brain. front. hum. neurosci. 3, 31 (2009), 2009. 2

[13] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 831–839, 2019. 2

[14] Ronald Kemker and Christopher Kanan. Fearnet: Brain-inspired model for incremental learning. *arXiv preprint arXiv:1711.10563*, 2017. 2

[15] Ronald Kemker and Christopher Kanan. Fearnet: Brain-inspired model for incremental learning. *arXiv preprint arXiv:1711.10563*, 2017. 3

[16] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5

[17] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017. 1

[18] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017. 2, 4, 6, 7

[19] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 4

[20] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 4

[21] Yuanpeng Li, Liang Zhao, Kenneth Church, and Mohamed Elhoseiny. Compositional language continual learning. In *International Conference on Learning Representations*, 2019. 3, 4, 5, 6

[22] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. 1

[23] Xialei Liu, Marc Masana, Luis Herranz, Joost Van de Weijer, Antonio M Lopez, and Andrew D Bagdanov. Rotate your networks: Better weight consolidation and less catastrophic forgetting. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 2262–2268. IEEE, 2018. 1, 2

[24] David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. *arXiv preprint arXiv:1706.08840*, 2017. 2

[25] David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. *arXiv preprint arXiv:1706.08840*, 2017. 3, 4, 7

[26] Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7765–7773, 2018. 2

[27] Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7765–7773, 2018. 2

[28] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier, 1989. 1

[29] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011. 4

[30] Cuong V Nguyen, Yingzhen Li, Thang D Bui, and Richard E Turner. Variational continual learning. *arXiv preprint arXiv:1710.10628*, 2017. 4, 6, 7

[31] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017. 2

[32] Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. Learning to learn without forgetting by maximizing transfer and minimizing interference. *arXiv preprint arXiv:1810.11910*, 2018. 3

[33] Anthony Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 7(2):123–146, 1995. 2

[34] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016. 1

[35] Andrei A Rusu, Neil C Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016. 2, 4, 5, 6

[36] Joan Serra, Didac Suris, Marius Miron, and Alexandros Karatzoglou. Overcoming catastrophic forgetting with hard attention to the task. In *International Conference on Machine Learning*, pages 4548–4557. PMLR, 2018. 2, 4, 5, 6, 7

[37] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. *arXiv preprint arXiv:1705.08690*, 2017. 2, 3

[38] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. *arXiv preprint arXiv:1705.08690*, 2017. 2

[39] Aaron Shon, Keith Grochow, Aaron Hertzmann, and Rajesh PN Rao. Learning shared latent structure for image synthesis and robotic imitation. In *Advances in neural information processing systems*, pages 1233–1240, 2006. 3

[40] Chenshen Wu, Luis Herranz, Xialei Liu, Yaxing Wang, Joost Van de Weijer, and Bogdan Raducanu. Memory replay gans: learning to generate images from new categories without forgetting. *arXiv preprint arXiv:1809.02058*, 2018. 2, 3

[41] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 374–382, 2019. 2

[42] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017. 4

[43] Chang Xu, Dacheng Tao, and Chao Xu. A survey on multi-view learning. *arXiv preprint arXiv:1304.5634*, 2013. 3

[44] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. *arXiv preprint arXiv:1708.01547*, 2017. 1, 2

[45] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. *arXiv preprint arXiv:1708.01547*, 2017. 2

[46] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pages 3987–3995. PMLR, 2017. 1, 2

[47] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pages 3987–3995. PMLR, 2017. 3

[48] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning*, pages 3987–3995. PMLR, 2017. 4

[49] Mengmi Zhang, Tao Wang, Joo Hwee Lim, Gabriel Kreiman, and Jiashi Feng. Variational prototype replays for continual learning. *arXiv preprint arXiv:1905.09447*, 2019. 4